



Europäisches Patentamt

European Patent Office

Office européen des brevets

(19)



EP 1 263 176 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

04.12.2002 Bulletin 2002/49

(21) Application number: 02019795.0

(22) Date of filing: 31.12.1999

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LU
MC NL PT SE

(30) Priority: 08.01.1999 EP 99100274

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:

99985683.0 7/1 142 226

(71) Applicant: Telefonaktiebolaget L M Ericsson
(Publ)

126 25 Stockholm (SE)

(54) Detection method and device

(57) In a data unit oriented communication between a sender and a receiver in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender. Methods and devices are described for detecting whether an acknowledgment data unit received by said sender corresponds to an original transmission or a retransmission.

(11)

(51) Int Cl: H04L 12/56, H04L 29/06

(72) Inventors:

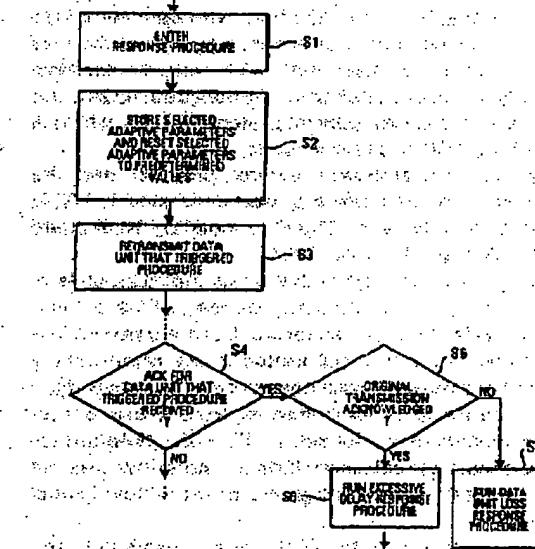
- Ludwig, Rainer
52383 Huerigenwald (DE)
- Meyer, Michael
52080 Aachen (DE)

(74) Representative: HOFFMANN - EITZ
Patent- und Rechtsanwälte
Arabellastrasse 4
81925 München (DE)

Remarks:

This application was filed on 05.09.2002 as a divisional application to the application mentioned under INID code 62.

FIG.1



EP 1 263 176 A2

EP 1 263 176 A2

Description

[0001] The present invention relates to detection methods and devices for use in a data unit oriented communication between a sender and a receiver in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender.

[0002] Data unit oriented communication is well-known. In data unit oriented communication an amount of data is divided into one or more data units, where the structure of the data units is defined by a communication protocol to which the sender and receiver in the communication adhere. The protocol also defines how specific information is to be coded, and how the sender and/or receiver may react to specific information. Data unit oriented communication is also known as packet exchange communication. It should be noted that the data units used in connection with specific protocols have different names, such as packets, frames, segments etc. For the purpose of the present description, the term "data unit" shall generically refer to all types of units used in a data unit oriented communication.

[0003] A feature that many communication protocols use for increasing reliability is that of acknowledging received data. More specifically, a sender or sending peer of the given protocol sends out data units, and the receiver or receiving peer of the given protocol acknowledges the correct receipt by returning appropriate acknowledgment data units. In this way, the sending peer is informed that the data units that were sent were also correctly received, and can accordingly adjust the flow control of the further data units to be sent. An example of a protocol that uses acknowledgment data units is the so-called transmission control protocol (TCP), which is a part of the TCP/IP protocol suite.

[0004] The transmission control protocol and the TCP/IP protocol suite are e.g. well described in "TCP/IP Illustrated, Volume 1 - The Protocols" by W. Richard Stevens, Addison-Wesley, 1994.

[0005] In order to cope with the fact that data units or acknowledgment data units may be lost, a time-out feature is provided in many protocols. Such a time-out feature means that a time-out period is set when data is sent, and if the specific data has not been acknowledged by the time the time-out period expires, a time-out response procedure is started. In TCP, the time-out response consists in retransmitting the data that was not acknowledged, and resetting one or more flow control parameters.

[0006] As an example, TCP uses a window-based flow control. TCP is a byte oriented protocol that divides a given number of bytes to be sent into so-called segments, and a record of the sent data is kept in terms of

bytes, i.e. up to which byte the data was sent, and a record of the received data is also kept in terms of bytes, i.e. up to which byte the data was received. The simplest way of controlling the flow of segments in connection with acknowledgment messages would be to send a segment and not send the next segment until the segment last sent was acknowledged. Such a method of flow control would however not be very efficient. As already mentioned, TCP uses window-based flow control, which is also referred to as flow control according to sliding windows. This concept is also well described in the above mentioned book by W. Richard Stevens.

[0007] Fig. 2 illustrates the concept of sliding windows. As can be seen, an amount of 8.192 bytes is to be sent in the example, where this amount is divided into 8 segments. The sending of segments is controlled in accordance with the send window, where the left end of the send window is defined by the data in the segments that have been sent and already acknowledged. In the example of Fig. 2 this is the data up to 2.048 bytes, i.e. the segments 1 and 2. The adjustment of the length of the send window, and thereby the right end of the window is a matter of the control procedure, which need not be explained in detail here.

[0008] The send window defines the amount of data which may have its corresponding acknowledgment outstanding. In the example of Fig. 2, the data up to 4.096 bytes, i.e. segments 3 and 4 have been sent and not yet acknowledged, and the difference between such sent and not acknowledged segments and the right end of the send window defines the usable window, i.e. the data that may still be sent without having received any further acknowledgments. As a consequence, in the example of Fig. 2 segments 5 and 6 may still be sent, but segments 7 and 8 can only be sent if the window moves to the right, which happens if further segments are acknowledged such that the left end moves to the right and/or if the length of the send window increases.

[0009] Furthermore, it should be noted that TCP provides for cumulative acknowledgment, i.e. there is not a one-to-one correspondence between segments and acknowledgments for segments, because one acknowledgment message may cover a plurality of segments. As an example, the receiving peer for the data amount shown in Fig. 2 could send an acknowledgment of bytes up to 4.096, such that this acknowledgment message would cover both segments 3 and 4.

[0010] The send window used by the sending peer will typically be determined by the so-called offered or advertised window, which is a data length provided to the sending peer by the receiving peer. In this way, the receiving peer can influence how many segments the sending peer will send at a time, and typically the advertised window will be calculated on the basis of the receiving peer's receive buffer. Also, the advertised window is a dynamic parameter that may be changed with every acknowledgment sent by the receiving peer.

[0011] Beyond the advertised window, it is also known

to define the so-called congestion window, which is used in connection with several congestion control routines such as slow start, congestion avoidance, fast retransmit and fast recovery," again see e.g. the above-mentioned book by W. Richard Stevens. The congestion window is a record that the sending peer keeps, and it is intended to take into account the congestion along the connection between the sending peer and receiving peer. As a typical control mechanism, the send window will be defined as the smaller of the advertised window and congestion window. [0012]. While the advertised window is a flow control imposed by the receiving peer, the congestion window is a flow control imposed by the sending peer, as a mechanism for taking congestion into account. [0013]. In a general sense, the congestion window is an example of an adaptive flow control parameter. In TCP the above-mentioned time-out response consists in resetting the congestion window to one segment and then consequently only sending one segment, namely retransmitting the segment that was not acknowledged and thereby caused the time-out. The sending peer then waits for the acknowledgment of said retransmitted segment. [0014]. Another example of an adaptive flow control parameter is the time-out period itself, which e.g. in TCP is referred to as RTO (Retransmission Time Out). The RTO is doubled as a response to a time-out. [0015]. As already mentioned, the time-out feature is a data loss detection mechanism. Other data loss detection mechanisms exist. Another example is the retransmission of data units in TCP in response to the receipt of duplicate acknowledgments. This mechanism will be briefly explained in the following. [0016]. As already mentioned (see e.g. Fig. 2), a data amount to be sent is divided into a sequence. Conventional implementations of TCP are arranged such that if the receiving peer has received and acknowledged a certain data amount up to a given byte (a certain number of consecutive segments), it expects the data that is next in the sequence. For example, if segments up to segment 4 have been received then segment 4 is acknowledged and the receiving peer expects to receive segment 5. If it then receives a further data unit that is different from segment 5 (e.g. segments 6, 7 and 8), it continues to acknowledge segment 4 for each data unit it receives. As a consequence, the sending peer receives duplicate acknowledgments. Commonly, TCP is implemented in such a way that the sending peer will count the number of duplicate acknowledgments, and if a certain threshold value is reached (e.g. 3), then the data unit next in the sequence to the data unit for which duplicate acknowledgments were received is retransmitted. [0017]. It is the object of the present invention to improve the communication in a system using a communications protocol that specifies the acknowledgment of sent data and specifies a data loss detection function.

such as a time-out function or a duplicate acknowledgment response function. [0018]. This object is achieved by detection methods and devices that are suitable for detecting whether an acknowledgment data unit received by a sender corresponds to an original transmission or a retransmission, as described in claims 1, 5, 9 and 11. [0019]. In accordance with an application of the detection method of the present invention, a sender in a communication will conduct a response procedure in response to an event that triggers a data loss detection mechanism, where the response procedure comprises at least two different modes for adapting the adaptive parameters used in flow control. In this way the method and device of the present invention are highly flexible in their management of triggering events, and can especially be implemented in such a way that the response procedure may be chosen depending on various potential causes of the triggering event, such that the correct responsive measures to a given situation may be invoked, and thereby measures can be avoided that might actually aggravate situations that may occur after a data loss detection mechanism was triggered. [0020]. The data loss detection mechanism is a mechanism that is capable of detecting a data loss. Examples are a time-out mechanism or a duplicate acknowledgment mechanism. Naturally, the invention may be applied to any suitable data loss detection mechanism. [0021]. A response procedure may comprise at least two different modes for adapting the adaptive parameters used in flow control. As an example, there are two modes, which are respectively associated with different causes of a time-out or a predetermined number of duplicate acknowledgments (e.g. the above-mentioned 3). More specifically, a first mode is associated with the loss of a data unit, and the second mode is associated with an excessive delay along the connection. Due to the use of two different modes, it is possible to adapt the parameters as is appropriate for the cause of the time-out or duplicate acknowledgments. Accordingly, the flow control procedure will contain one or more evaluation and judgment steps, in which the triggering event is qualified, e.g. a categorization is conducted as to what caused the event. Then, depending on the result of this characterization, an appropriate response procedure may be enabled. In the context of the above example, if it is determined that the time-out or duplicate acknowledgments are caused by the loss of a data unit, then the known response procedure to the loss of data units may be run, as it is, e.g. known from conventional TCP, which assumes that any time-out or the receipt of several duplicate acknowledgments is caused by the loss of a data unit. There is, however, a second mode, and if it is determined that the time-out or duplicate acknowledgments are caused by an excessive delay along the connection, then an excessive delay response procedure is run, which will typically be different from the response procedure to the loss of a data unit.

[0022] More specifically, as will also be explained in more detail in the following, the judgment that data units have been lost will be answered by reducing the transmission rate to thereby avoid further congestion. On the other hand, if there is excessive delay along the connection, then the measures taken in response to a supposed loss of data units would not be helpful, much rather they might actually aggravate the problem causing the excessive delay. Consequently, the response procedure to excessive delay will typically be different, and e.g. comprise keeping the transmission rate at the previous level, but on the other hand increasing the time-out period, such that further unnecessary retransmissions are avoided.

[0023] Naturally, the an application of the detection method of the present invention may be implemented as providing an arbitrary number of modes or response procedures to various causes of triggering events. The number of modes and the specific measures taken in each mode naturally depend on the specific situation, i.e. the chosen protocol, the given communication situation, etc.

[0024] An important aspect is that although the data loss detection mechanism is capable of detecting data loss, the reaction to the triggering of the data loss detection mechanism does not assume that a data loss has necessarily occurred, much rather a 'flexible' response is possible, which may take into account various causes of the triggering event.

[0025] Further aspects and advantages of the present invention shall be better understood from the following detailed description, which makes reference to the figures, in which:

Fig. 1 shows an embodiment of an overall control procedure applying the detection method of the present invention;

Fig. 2 is an explanatory diagram for describing the concept of window-based flow control;

Fig. 3 is a graph for explaining the advantages of the present invention; and

Fig. 4 is an explanatory diagram for illustrating a situation in which an excessive delay may be caused in a connection between two host computers.

[0026] Although the following description will be generally directed towards any communications protocol that makes use of data acknowledgement and also provides a time-out feature, examples will often be given that relate to the transmission control protocol TCP known from the TCP/IP protocol suite. The application of the present invention to this protocol is a preferred embodiment. In order to avoid any unnecessary repetition, the disclosure in the introduction of this application

is incorporated into the invention disclosure.

[0027] Fig. 1 shows a partial flow diagram. As can be seen, step S1 indicates that a response procedure is entered. Fig. 1 does not show the flow control procedure leading up to this point, as it is of no importance for the present explanation. For example, it may be the window-based flow control procedure explained in connection with Fig. 2 and e.g. well known from TCP. It is only important for the invention that there is data acknowledgement and a data loss detection feature, such that a sending peer of the protocol has the capability of detecting a possible or potential data loss, and may conduct a corresponding response procedure. As already mentioned, the data loss detection feature may e.g. be a time-out feature or a duplicate acknowledgment detection feature.

[0028] In the example of Fig. 1, after the response procedure is entered, selected adaptive parameters that are used for the flow control are stored and then reset to predetermined values in step S2. As an example, the time-out period and/or the above described congestion window are such adaptive flow control parameters. In conventional TCP, the congestion window is typically reset to a value of one segment and at the same time the RTO is doubled. It should be noted that not all adaptive parameters used in the flow control procedure need to in fact be changed, much rather only a selected number.

[0029] Also, it should be clear that the present invention is naturally not restricted to window-based flow control and the associated adaptive parameters, much rather the invention is applicable to any flow control principle and the associated adaptive parameters.

[0030] Returning to Fig. 1, the data unit that triggered the event (e.g. caused a time-out) is retransmitted in step S3. In other words, when staying with the example of a time-out, the data unit for which no acknowledgment was received during the time-out period is retransmitted. Then, at a later point it is determined in step S4 if an acknowledgment associated with the retransmitted data unit has been received. This may be a cumulative acknowledgment or also a single acknowledgment. It may be noted that the dotted lines in Fig. 1 indicate that other steps may be interposed, but these are of no importance to the present invention. Then, according to the example of Fig. 1, in accordance with the present invention step S5 determines if the acknowledgment associated with the data unit that was retransmitted in fact acknowledges the original transmission of the data unit or the retransmission. It should be noted that the "original transmission" may already be a retransmission, such that the "retransmission" may be the retransmission of a retransmission etc. There are various possibilities of implementing step S5, as will be explained further on.

[0031] If step S5 determines that the acknowledgment message in fact acknowledges the retransmission of the data unit, then the procedure goes to step S7, in which a data unit loss response procedure is run, because the negative outcome of the decision step S5 in-

dicates that the original transmission of the data unit was lost. In the example of TCP, step S7 will consist in conventional measures against data unit loss. [0032] On the contrary, if the decision step S5 is answered in the affirmative, then the procedure goes to step S6, in which a response procedure is run that answers an excessive delay. In other words, because step S5 indicated that in fact the original transmission of the data unit was not lost, but only excessively delayed, corresponding measures must be taken. For example, when taking TCP as a protocol example, this may consist in returning the congestion window to the value stored in step S2 and on the other hand adapting the time out period to the delay. In other words, the round trip time RTT associated with the original transmission and the acknowledgment of the original transmission can be used as a basis for adapting the time out period. Thereby, further unnecessary retransmissions and time-outs or duplicate acknowledgments due to excessive delay can be avoided. [0033] Preferably, the congestion window is not simply reset to the previous value, but much rather is set to the value it would have assumed, had the response procedure not taken place, i.e. had the data loss detection mechanism not been triggered. [0034] As can be seen, the example of Fig. 1 shows a first mode consisting of steps S2, S3, S4, S5, and S7, as well as a second mode consisting of steps S2, S3, S4, S5 and S6. [0035] In order to better explain the above example, reference will now be made to Fig. 3, which shows an example of a flow control procedure conducted in connection with conventional TCP. The graph shows the amount of data in bytes transported over time. As can be seen, the first two segments are sent at time t=4s. Then, due to the interaction of receiving acknowledgment data units and the adjustment of adaptive parameters not shown, segments are sent. [0036] For the purpose of explanation, it should be noted that the diamond shaped symbols refer to segments, and the square symbols to acknowledgment data units. The diamond symbols indicate the first byte of the segment, whereas the squares indicate the lowest unacknowledged byte. The acknowledgment data units indicated at a certain segment level always acknowledge the sent segments up to that segment level. In other words, the acknowledgment at a segment level of 8.400 bytes (t=12s) acknowledges the segments below 8.400 byte, but not including byte 8.400. Quite to the contrary, as explicitly indicated in the graph, the segment at 6.400 byte (t=10s) is a data unit or packet that causes a time-out. As a consequence, a retransmission is conducted of said data unit at the 6.400 byte level. [0037] Now, if it is assumed that the time-out shown in Fig. 3 was caused by an excessive delay and not by the shown first packet being lost, then the retransmission has the following negative consequences. [0038] For one thing, it leads to a decreased through-

put performance, as the same data has to traverse the connection or connecting path twice, which wastes bandwidths that could have otherwise been used for useful data. This negative consequence will occur in any protocol that falsely responds to a time-out by retransmitting the data unit. [0039] If, as shown in Fig. 3, the TCP protocol is used, then the reaction of the sending peer to such a time-out not caused by data unit loss is particularly disadvantageous: the sender will retransmit all outstanding packets and above that reduce its transmission rate. This is explicitly shown in Fig. 3. [0040] It may be noted that the above described time-out not caused by data unit loss is also referred to as a spurious time-out. [0041] As also shown in Fig. 3, in conventional TCP the sender misinterprets all acknowledgments associated with retransmitted data units as acknowledging the retransmission, even though these acknowledgments (ACKs) in fact are delayed acknowledgments of the original transmissions. [0042] What Fig. 3 does not show, is that additionally the duplicate data units sent by the sending peer will trigger duplicate acknowledgments at the receiving peer, which will lead to yet another reduction in the transmission rate at the conventional TCP sender, namely the congestion window is set to one half of its earlier value. [0043] The occurrence of excessive delay that goes beyond what the TCP time-out period can account for, may especially appear in wireless networks or such protocol connections of which at least a part runs over a wireless link. The inventors of the present application realized that spurious time-outs can happen often enough in such networks, so that serious performance degradation results. Examples of this will now briefly be mentioned. [0044] Fig. 4 shows a situation, where two host computers act as peers of the TCP (indicated by the long arrows from host to host at the bottom and top of the figure). The lower protocol layers comprise a radio link over a wireless access network to the internet. The connection between the internet and the host on the right is not shown. An example of a protocol for the radio link is the so-called radio link control protocol RLC. As indicated in Fig. 4, both the transport layer protocol (e.g. TCP) and the link layer protocol (e.g. RLC) have an ARQ (Automatic Retransmission reQuest) function. This means that these protocols both implement time-out and retransmission functions. In the situation of Fig. 4, due to the ARQ being used at the link layer, a race condition is generated between the link layer and the transport layer: while the link layer retransmits data, the transport retransmission timer might expire, leading to a spurious time-out. The retransmissions at the link layer can be due e.g. to transmission errors or to data loss because of handovers. [0045] It may also be noted that the transmission de-

lay over the wireless network is often a considerable fraction of the end-to-end delay between the sending and receiving peer of the transport layer protocol. If in this case the bandwidth available to the transport layer connection in the wireless network drops considerably over a short period of time, the resulting increase in the end-to-end delay between the transport layer sender and receiver might lead to spurious time-outs. Examples of bandwidth drops include mobile hosts executing a handover into a cell which provides less bandwidth than the old cell. [0046] As already indicated previously, when employing the present invention, the problem described in connection with Fig. 3 can be avoided. More specifically, when applying the method described in connection with Fig. 1 to the problem in Fig. 3, then the sending peer is capable of distinguishing between acknowledgment data units to the original transmission of a data unit and acknowledgment data units to the retransmission of a data unit. From this information, the sender can decide if a spurious time-out has occurred, or if there indeed has been a loss of a data unit. The sender can then react accordingly.

[0047] More specifically, in the example of Fig. 3, the sender using the invention will be able to identify the acknowledgment data unit received after having retransmitted the shown first packet as being an acknowledgment for the original transmission ($t=10s$) and not for the retransmission ($t=15s$). Due to this, the sender will perform an appropriate response procedure to the excessive delay, namely not retransmit the data units following the first retransmitted data unit, and also not decrease the transmission rate, much rather the sender will increase the time-out period employed in the flow control on the basis of the measured delay between the original sending of the data unit and the receipt of the corresponding acknowledgment data unit for said original sending. In this way, further spurious retransmissions and time-outs can be avoided.

[0048] As may be seen, the present invention is capable of providing a mechanism that allows a more flexible communication system when using a protocol that provides acknowledgment of data and a time-out function or duplicate acknowledgment detection function. In the example just described, the application of the invention is capable of qualifying a triggering event, i.e. distinguishing between at least two different causes, and then capable of invoking an appropriate response procedure. It may be noted that in the above examples the modes for adapting the adaptive parameters were associated with data unit loss on the one hand and excessive delay on the other, but naturally the application of the present invention is by no means restricted thereto. Much rather, the modes for adapting the adaptive parameters may be associated with any possible cause of time-out events or duplicate acknowledgment events.

[0049] In the example described in Fig. 1, it was decided in step S6, which implements the present invention, if the acknowledgment data unit associated with a given data unit acknowledged the original transmission or the retransmission of said given data unit. According to a first preferred embodiment for implementing this step, the sender keeps a record of the round trip time RTT associated with the connection between sending and receiving peer, and especially keeps a record of the shortest RTT found during the connection or session up to the point of time under consideration. Then, if an acknowledgment data unit for a retransmitted data unit is received within a time period that is smaller than a predetermined fraction of said shortest RTT, then the sender determines that this acknowledgment belongs to the original transmission and not the retransmission. This fraction may be set to a fixed value, or may itself be an adaptive parameter. Naturally, it is not necessary that the comparison value multiplied with said fraction is the shortest measured RTT, much rather it is also possible that the sender keeps an average RTT value. In this sense, the comparison value to be multiplied by said fraction is generally a function of one or more RTT values measured in the course of the connection (during the session).

[0050] According to another preferred embodiment for implementing step S5, the sender adds a mark to data units that it sends, where said mark is defined in such a way that it allows to distinguish between an original transmission and a retransmission. Then, the receiver can accordingly mark acknowledgment data units, such that the sender is capable of identifying if an acknowledgment refers to the original transmission or the retransmission.

[0051] This marking of data units can be done in any desired way. For example, it would in theory be possible to simply designate a single bit in the data units, where a value of 0 would indicate original transmission and a value of 1 a retransmission, *vice versa*. In a general sense, a bit string can be chosen that may also convey some more information. However, in connection with protocols that provide for such an option, it is preferred to use the time stamp option. This option is e.g. well known for TCP, see the above mentioned book by W. R. Stevens. In other words, it is preferred to include a time stamp in sent data units, which indicates when the data unit was sent. The receiver can then simply include the same time stamp in the acknowledgment data unit, so that the sender has a unique way of identifying the data units to which the acknowledgment refers.

[0052] Further overall method and system embodiments:

E1. A method for controlling a data unit oriented communication between a sender and a receiver operating in accordance with a predetermined communication protocol, where
said sender divides an amount of data to be sent into one or more data units having a structure determined by said protocol,

said receiver acknowledges the correct receipt of data units by returning acknowledgment data units to said sender, wherein said data units are sent by said sender in accordance with a flow control procedure conducted on the basis of one or more adaptive parameters and said acknowledgment data units, and wherein said flow control procedure comprises a data loss detection mechanism capable of detecting data loss in said communication, said data loss detection mechanism being triggered to indicate the potential loss of data by one or more predetermined events, where in response to the triggering of said data loss detection mechanism a corresponding response procedure is conducted, said response procedure comprising at least two different modes for adapting said one or more adaptive parameters.

E2. A method according to embodiment E1, wherein said data loss detection mechanism is a time out mechanism, such that after a given data unit is sent, said sender monitors a time out period and if no acknowledgment data unit associated with said given data unit is received before said time out period expires, said time out mechanism is triggered.

E3. A method according to embodiment E1, wherein said data loss detection mechanism is a duplicate acknowledgment detection mechanism, such that said sender monitors the received acknowledgments, and if a data unit is acknowledged a predetermined number of times, said duplicate acknowledgment detection mechanism is triggered.

E4. A method according to embodiment E2 or E3, wherein said response procedure comprises the retransmission of said given data unit.

E5. A method according to embodiment E4, wherein the decision on which of said at least two modes to choose for adapting said adaptive parameters is made on the basis of one or more acknowledgment data units received by said sender after having retransmitted said given data unit.

E6. A method according to embodiment E2, wherein said time out period is one of said adaptive parameters.

E7. A method according to one of embodiments E1 to E6, wherein said flow control procedure is window based, and one or more flow control windows are among said adaptive parameters.

E8. A method according to embodiment E5, wherein said at least two modes consist of a first and a second mode, said first mode being associated with the judgment that the triggering event was caused

by the loss of said given data unit, and said second mode being associated with the judgment that said given data unit or the acknowledgment data unit for said given data unit has been excessively delayed.

E9. A method according to embodiment E8, wherein said sender marks data units being sent such that an original transmission may be distinguished from a retransmission, and said receiver correspondingly marks the acknowledgment data units, such that the acknowledgment of an originally sent data unit may be distinguished from the acknowledgment of the retransmission of said data unit.

E10. A method according to embodiment E9, wherein the sender marks data units by including a time stamp in each sent data unit, said time stamp indicating the time said data unit was sent, and the receiver marks the acknowledgment data unit for a received data unit by including the time stamp contained in said received data unit in the acknowledgment data unit for said received data unit.

E11. A method according to embodiment E9, wherein the sender marks data units by including a bit string in each sent data unit, said bit string having at least two different values for distinguishing between an original transmission and a retransmission, and the receiver marks the acknowledgment data unit for a received data unit by including the bit string contained in said received data unit in the acknowledgment data unit for said received data unit.

E12. A method according to embodiment E11, wherein said bit string consists of a single bit.

E13. A method according to embodiment E11 or E12, wherein said first mode is chosen if the first acknowledgment data unit associated with said given data unit that is received after having retransmitted said given data unit acknowledges the retransmission of said given data unit, and said second mode is chosen if the first acknowledgment data unit associated with said given data unit that is received after having retransmitted said given data unit acknowledges the original transmission of said given data unit.

E14. A method according to embodiment E9, wherein the sender measures the round trip time associated with the connection for sending of said amount of data, and the receiver measures the round trip time associated with the connection for receiving of said amount of data.

the time between the retransmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit is determined and compared to a value derived from one or more of said round trip time measurements, and said first or second mode is chosen on the basis of the result of said comparison.

E15. A method according to embodiment E14, wherein said value derived from said round trip time measurements is the shortest round trip time for the connection, and the second mode is chosen if said time between the retransmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit is smaller than a predetermined fraction of said smallest round trip time.

E16. A method according to one of embodiments E9 to E15, wherein the second mode comprises adapting the time out period on the basis of the time that elapsed between the original transmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit.

E17. A method according to one of embodiments E9 to E16, wherein the flow control procedure is window based and a congestion window is used, where the value of said congestion window at the time of said triggering event is stored after said triggering event occurred and subsequently said value of the congestion window is reset to a predetermined value, and if said second mode is chosen after having received the first acknowledgment data unit associated with said given data unit, said value of said congestion window is set to the value it would have assumed, had the response procedure not taken place.

E18. A communication device for data unit oriented communication in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units having a structure determined by said protocol and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender, and where said communication device, when acting as a sender, is arranged to send data units in accordance with a flow control procedure that is conducted on the basis of one or more adaptive parameters and said acknowledgment data units, and said flow control procedure comprising a data loss detection mechanism capable of detecting data

loss in said communication, said data loss detection mechanism being triggered to indicate the potential loss of data by one or more predetermined events, where in response to the triggering of said data loss detection mechanism a corresponding response procedure is conducted, said response procedure comprising at least two different modes for adapting said one or more adaptive parameters.

E19. A device according to embodiment E18, wherein said data loss detection mechanism is a time out mechanism, such that after a given data unit is sent, said sender monitors a time out period, and if no acknowledgment data unit associated with said given data unit is received before said time out period expires, said time out mechanism is triggered.

E20. A device according to embodiment E18, wherein said data loss detection mechanism is a duplicate acknowledgment detection mechanism, such that said sender monitors the received acknowledgments, and if a data unit is acknowledged a predetermined number of times, said duplicate acknowledgment detection mechanism is triggered.

E21. A device according to embodiment E20, wherein the device is arranged to act as a sender and as a receiver, said device when acting as a sender is arranged to divide said amount of data into a defined sequence, said device when acting as a receiver is arranged to monitor the order of received data units in terms of said sequence, such that if data units are correctly received that are out of order, the last data unit correctly received is acknowledged in order of receipt of a data unit out of order, and said duplicate acknowledgment detection mechanism is such that after a given data unit is sent, said sender monitors the received acknowledgments, and if the data unit preceding said given data unit is acknowledged a predetermined number of times, said duplicate acknowledgment detection mechanism is triggered.

E22. A device according to embodiment E19 or E21, wherein said response procedure comprises the retransmission of said given data unit.

E23. A device according to embodiment E22, wherein the decision on which of said at least two modes to choose for adapting said adaptive parameters is made on the basis of one or more acknowledgment data units received by said sender after having retransmitted said given data unit.

E24. A device according to embodiment E19, wherein said time out period is one of said adaptive parameters.

E25. A device according to one of embodiments E18 to E24, wherein said flow control procedure is window based, and one or more flow control windows are among said adaptive parameters.

E26. A device according to embodiment E23, wherein said at least two modes consist of a first and a second mode, said first mode being associated with the judgment that the triggering event was caused by the loss of said given data unit, and said second mode being associated with the judgment that said given data unit or the acknowledgment data unit for said given data unit has been excessively delayed.

E27. A device according to embodiment E26, wherein said device when acting as a sender marks data units being sent such that an original transmission may be distinguished from a retransmission, and said device when acting as a receiver correspondingly marks the acknowledgment data units, such that the acknowledgment of an originally sent data unit may be distinguished from the acknowledgment of the retransmission of said data unit.

E28. A device according to embodiment E27, wherein the device when acting as a sender marks data units by including a time stamp in each sent data unit, said time stamp indicating the time said data unit was sent, and the device when acting as a receiver marks the acknowledgment data unit for a received data unit by including the time stamp contained in said received data unit in the acknowledgment data unit for said received data unit.

E29. A device according to embodiment E27, wherein the device when acting as a sender marks data units by including a bit string in each sent data unit, said bit string having at least two different values for distinguishing between an original transmission and a retransmission, and the device when acting as a receiver marks the acknowledgment data unit for a received data unit by including the bit string contained in said received data unit in the acknowledgment data unit for said received data unit.

E30. A device according to embodiment E28 or E29, wherein said first mode is chosen if the first acknowledgment data unit associated with said given data unit that is received after having retransmitted said given data unit acknowledges the retransmission of said given data unit, and said second mode is chosen if the first acknowledgment data unit associated with said given data unit that is received after having retransmitted said given data unit acknowledges the original transmission of said given data unit.

E31. A device according to embodiment E28, and wherein the device when acting as a sender measures the round trip time associated with the connection for sending of said amount of data, the time between the retransmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit is determined and compared to a value derived from one or more of said round trip time measurements, and said first or second mode is chosen on the basis of the result of said comparison.

E32. A device according to embodiment E31, wherein said value derived from said round trip time measurements is the shortest round trip time for the connection, and the second mode is chosen if said time between the retransmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit is smaller than a predetermined fraction of said smallest round trip time.

E33. A device according to one of embodiments E26 to E32, wherein the second mode comprises adapting the time out period on the basis of the time that elapsed between the original transmission of said given data unit and the receipt of the first acknowledgment data unit associated with said given data unit.

E34. A device according to one of embodiment E28 to E33, wherein the flow control procedure is window based and a congestion window is used, where the value of said congestion window at the time of said triggering event is stored after said triggering event occurred and subsequently said value of the congestion window is reset to a predetermined value, and if said second mode is chosen after having received the first acknowledgment data unit associated with said given data unit, said value of said congestion window is set to the value it would have assumed, had the response procedure not taken place.

Claims

1. In a data unit oriented communication between a sender and a receiver, in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender.

1. A method of detecting whether an acknowledgment data unit received by said sender corresponds to an original transmission or a retransmission, comprising:

5 said sender adding a mark to data units that it sends, said mark being defined in such a way that it allows to distinguish between an original transmission and a retransmission,

10 said receiver correspondingly adding the mark received in a data unit to an acknowledgment data unit for said received data unit, and,

15 said sender determining whether the received acknowledgment data unit is associated with an original transmission or a retransmission based on said mark.

2. The method of claim 1, wherein said mark is a predetermined bit string.

20

3. The method of claim 1 or 2, wherein said mark is a time stamp.

25

4. The method of one of claims 1 to 3, wherein said predetermined communication protocol is the Transmission Control Protocol.

30

5. A communication device for data unit oriented communication in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender, where said communication device is arranged such that when acting as a sender, it adds a mark to data units that it sends, said mark being defined in such a way that it allows to distinguish between an original transmission and a retransmission, and determines whether a received acknowledgment data unit is associated with an original transmission or a retransmission, based on said mark.

35

6. The device of claim 5, wherein said mark is a predetermined bit string.

40

7. The device of claim 5 or 6, wherein said mark is a time stamp.

45

8. The device of one of claims 5 to 7, wherein said predetermined communication protocol is the Transmission Control Protocol.

50

9. In a data unit oriented communication between a sender and a receiver in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units by returning acknowledgment data units to the sender, a method of detecting whether an acknowledgment data unit received by said sender corresponds to an original transmission or a retransmission, comprising:

55 the sender measuring the round trip time associated with the connection between said sender and receiver,

60 determining the time between a retransmission of a data unit and the receipt of the first acknowledgment data unit associated with said data unit and comparing said time to a value derived from one or more of said round trip time measurements,

65 determining whether the received acknowledgment data unit is associated with an original transmission or a retransmission based on said comparison.

70

10. The method of claim 9, wherein said value derived from said round trip time measurements is the shortest round trip time for the connection, and said received acknowledgment data unit is determined to be associated with a retransmission if the time between the retransmission of said data unit and the receipt of the first acknowledgment data unit associated with said data unit is smaller than a predetermined fraction of said round trip time.

75

11. A communication device for data unit oriented communication in accordance with a predetermined communication protocol, where said communication protocol prescribes that the sender in a communication divides an amount of data to be sent into one or more data units, and the receiver in said communication acknowledges the correct receipt of data units to the sender, where said communication device is arranged such that when acting as a sender, it measures the round trip time associated with the connection between said sender and receiver, determines the time between a retransmission of a data unit and the receipt of the first acknowledgment data unit associated with said data unit and compares said time to a value derived from one or more of said round trip time measurements, and determines whether the received acknowledgment data unit is associated with an original transmission or a retransmission based

19.

EP 1 263 176 A2

20

on said comparison.

12. The device of claim 11, wherein said communication device is furthermore arranged such that said value derived from said round trip time measurements is the shortest round trip time for the connection, and said received acknowledgment data unit is determined to be associated with a retransmission if the time between the retransmission of said data unit and the receipt of the first acknowledgment data unit associated with said data unit is smaller than a predetermined fraction of said round trip time.

10

15

20

25

30

35

40

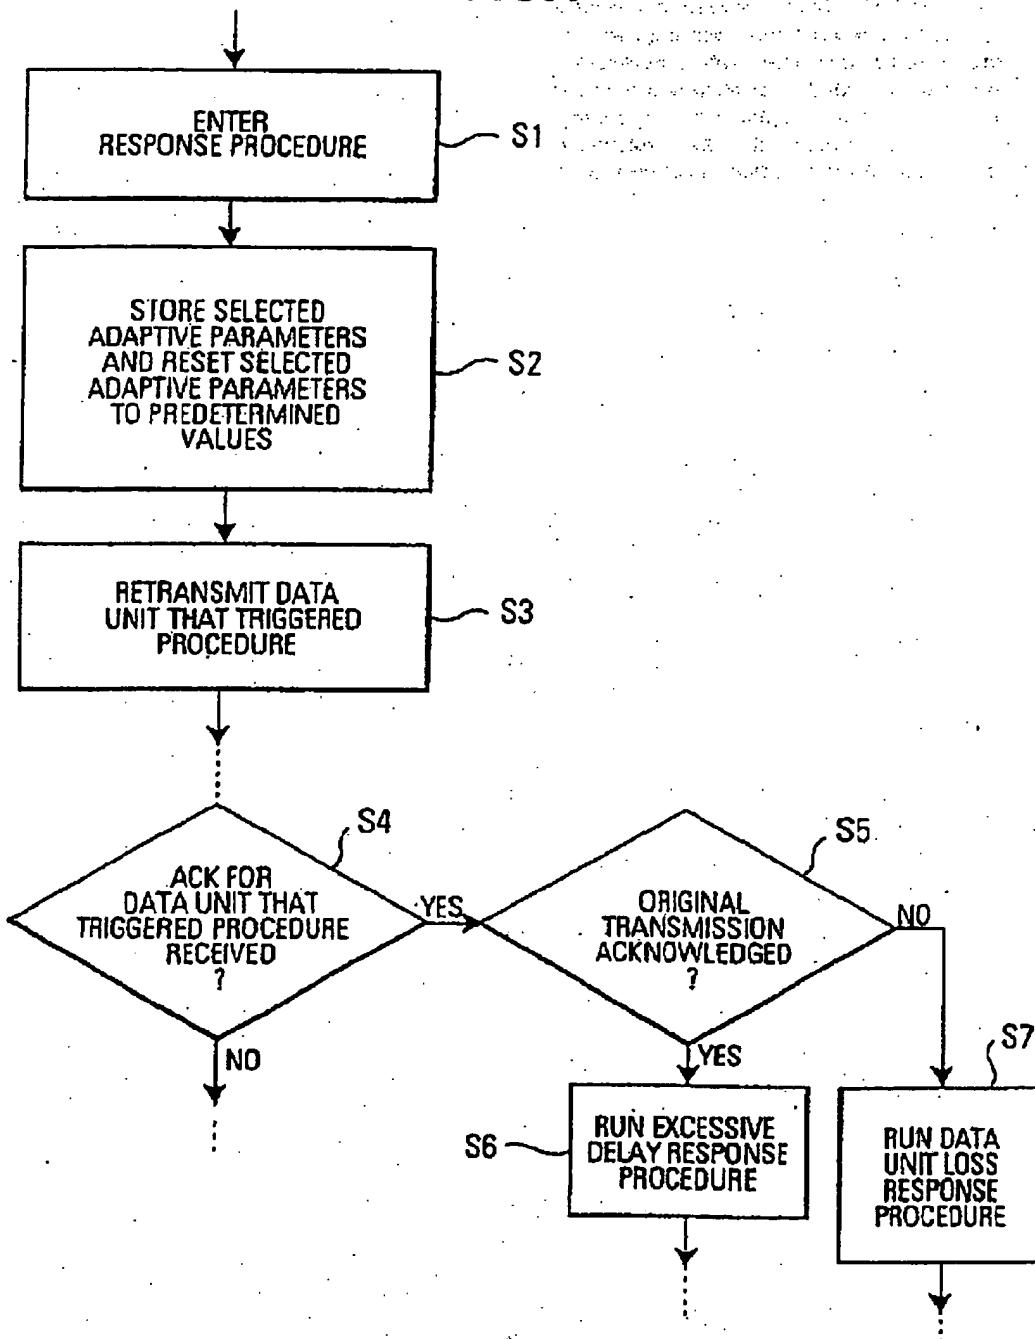
45

50

55

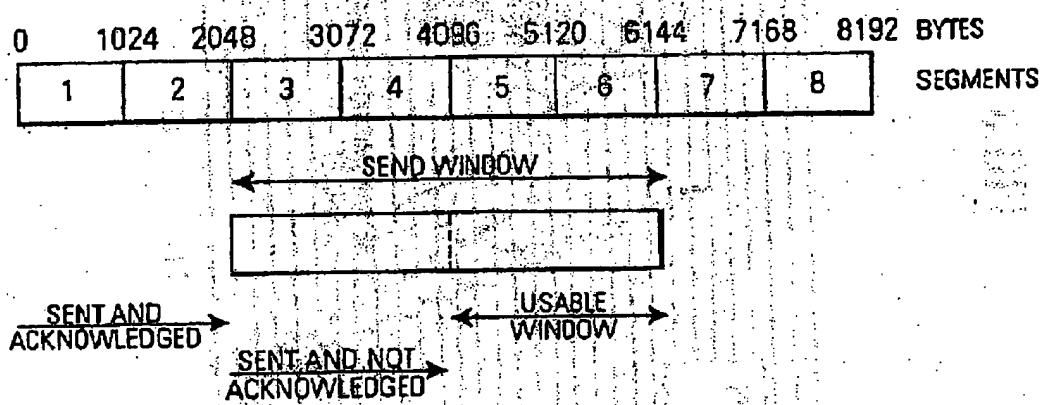
EP 1 263 176 A2

FIG.1



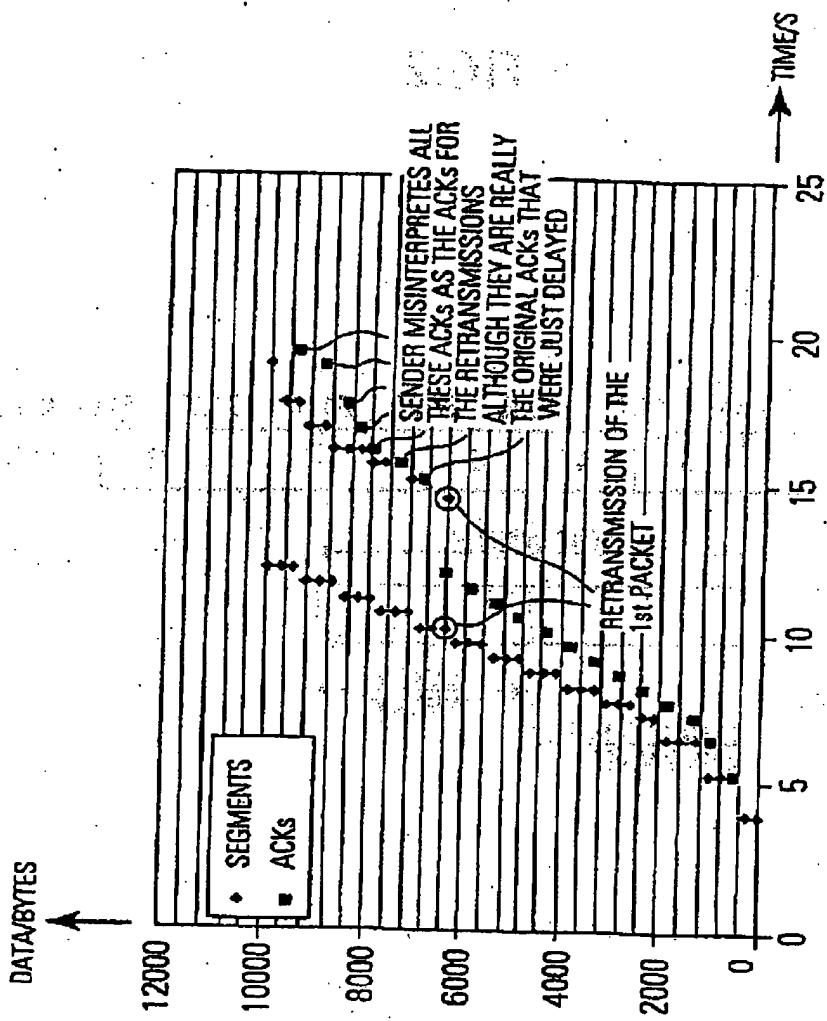
EP 1.263.176 A2

FIG.2



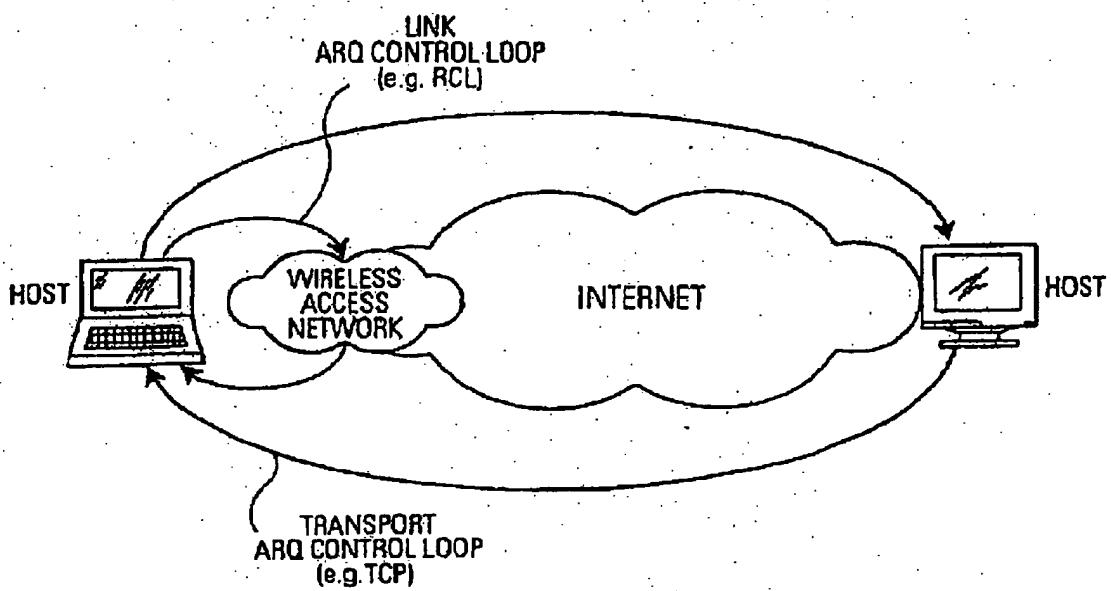
EP 1 263 176 A2

FIG.3



EP 1 263 176 A2

FIG.4



EP 1 263 176 A3

European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 02 01 9795

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (INCL7)
A	<p>BALAKRISHNAN H ET AL; "A COMPARISON OF MECHANISMS FOR IMPROVING TCP PERFORMANCE OVER WIRELESS LINKS" IEEE / ACM TRANSACTIONS ON NETWORKING, vol. 5, no. 6, 1 December 1997 (1997-12-01), pages 756-769, XP000734405 * page 756, left-hand column, line 1 - page 758, left-hand column, line 15 * * page 760, left-hand column, line 4 - line 22 * * page 767, right-hand column, line 52 - page 768, left-hand column, line 28 *</p>	1-12	H04L12/56 H04L29/06
A	<p>WO 98 37670 A (AT&T CORP) 27 August 1998 (1998-08-27) * page 2, line 15 - page 3, line 16 * * page 4, line 18 - line 33 * * page 5, line 21 - page 6, line 28 * * page 9, line 5 - page 12, line 6 * * page 14, line 12 - page 15, line 2 * * claims 1-6 *</p>	1-12	<p>TECHNICAL FIELDS SEARCHED (INCL7)</p> <p>H04L H04Q</p>
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of completion of the search		Examiner
THE HAGUE	13 January 2003		Karavassis, M
CATEGORY OF CITED DOCUMENTS			
X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background C: non-patent literature P: intermediate documents			
T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons S: member of the same patent family, corresponding document			

EP 1 263 176 A3

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 02 01 9795

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The numbers are as contained in the European Patent Office EPO file on. The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

13-01-2003

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9837670 A	27-08-1998 US WO	5974828 A 9837670 A1	26-10-1999 27-08-1998

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

8
CORONA PX

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.